The Effect of Temporal Aggregation on Battery Sizing for Peak Shaving

Dominik Werle[†]
dominik.werle@kit.edu

Simon Bischof[‡]
simon.bischof2@kit.edu

Holger Trittenbach[†]
holger.trittenbach@kit.edu

Daniel Warzel[†]
daniel.warzel@student.kit.edu

Anne Koziolek[†]
anne.koziolek@kit.edu

Klemens Böhm[†]
klemens.boehm@kit.edu

[†]Karlsruhe Institute of Technology (KIT), Germany
Institute for Program Structures and Data Organization (IPD)

[‡]Karlsruhe Institute of Technology (KIT), Germany
Institute for Data Processing and Electronics (IPE)

ABSTRACT

Battery systems can reduce the peak electrical consumption through proper charging and discharging strategies. To this end, consumers often rely on historic consumption data to select a cost-efficient battery system. However, historic data is an imperfect mapping of the real consumption, because of a coarse sampling rate or measurement inaccuracies. This can result in non-optimal decisions, e.g., by underestimating the battery capacity required. In this article, we analyze how aggregation affects a state-of-the-art battery sizing algorithm for an industrial production site. We then use machine learning on a short period of high-resolution data to correct this error from historic data. Our experiments indicate that machine learning models can correct this error in some cases. However, adding a safety margin obtained from historic data to the battery size is a more reliable way of reducing the error.

CCS CONCEPTS

• **Information systems** → **Data compression**; • **Applied computing** → *Electronics*; • **Hardware** → *Batteries*; *Power estimation and optimization.*

KEYWORDS

Storage Sizing, Data Compression, Time Series Aggregation

ACM Reference Format:

Dominik Werle, Simon Bischof, Holger Trittenbach, Daniel Warzel, Anne Koziolek, and Klemens Böhm. 2019. The Effect of Temporal Aggregation on Battery Sizing for Peak Shaving. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems (e-Energy '19), June 25–28, 2019, Phoenix, AZ, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3307772.3331023

e-Energy '19, June 25–28, 2019, Phoenix, AZ, USA

ACM ISBN 978-1-4503-6671-7/19/06…$15.00
https://doi.org/10.1145/3307772.3331023

1 INTRODUCTION

Battery systems enable shifting times when energy is sourced from the grid [7–9]. While there often are alternatives such as load shifting which do not require a battery system, this requires demand side flexibility. Such flexibility does not always exist [2].

Selecting a good configuration of a battery system is difficult. It involves selecting the physical and chemical make-up and the charging strategy [1]. When consumption does not change much over time, one can rely on historic data to this end [6].

However, historic data is an imperfect representation of the real consumption and subject to disturbances. Also, data often is aggregated to minimize storage and computing infrastructure requirements, resulting in loss of information for analyses [10]. So deciding on a battery configuration based on historic data can be non-optimal and cost inefficient. For example, when using batteries to reduce peak consumption, aggregated historic data may suggest battery sizes smaller than needed. This in turn is costly since large consumers often are billed based on peak consumption. Therefore, to make good investment decisions it is important for a system planner to get information (a) on the expected error of the battery sizing results, and (b) on the impact of using better measurement infrastructure or more fine-grained measurements.

Example 1. *Think of a plant manager who thinks about installing a battery system to reduce peak consumption. Suppose that there already is a large data set of historic consumption data available, collected at 15 minute resolution. The manager intends to use this data to decide on the size of the new battery. She also has the option to install a high-resolution measurement device to collect additional data for a certain period of time. However, it is unclear whether this additional data is useful to improve the investment decision.*

This scenario gives way to the following questions. $\mathbf{Q_1}$: How does data aggregation affect solutions of battery sizing optimization? $\mathbf{Q_2}$: Is it possible to estimate this effect from a short period of fine-grained measurements?

Contributions. In this article, we study these questions based on smart-meter measurements from an industrial production site [3]. We make two specific contributions. $\mathbf{C_1}$: We formulate an optimization problem for battery-based peak shaving, to optimize battery sizes for consumption time series and target maximum peak loads. Based on it, we analyze the influence of aggregation intervals and

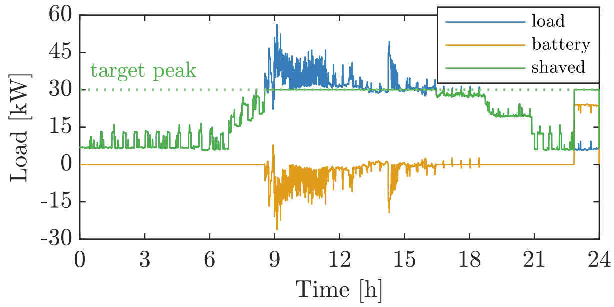

Figure 1: Example application of a 30 kWh battery for battery-pased peak shaving on a one day load profile.

aggregation types and of the target maximum peak consumption on the optimal battery size. $\mathbf{C_{2a}}$: We propose using a short period of high-resolution measurements as training data to predict the errors on battery sizing based on aggregated, historic data.[1] $\mathbf{C_{2b}}$: We apply machine learning to correct non-optimal battery sizes.

Results. A core takeaway from our article is that common aggregation levels of 15 min or more affect the battery sizing significantly. Lower aggregation levels help in making better sizing decisions. Particularly, aggregation leads to underestimated battery sizes. A second takeaway is that machine learning models trained on a short period of high-resolution data do not significantly reduce this error in our experiments. We conclude that adding a safety margin to the battery size obtained from historic data is a more reliable way of reducing the error for this use case.

2 BATTERY-BASED PEAK SHAVING

In this section, we introduce peak shaving with minimal battery size as an optimization problem. The idea is to find a minimum battery size so that the consumption from the grid does not exceed a given threshold over a selected time horizon. To illustrate, Figure 1 shows a daily load profile that is shaved to the target maximum peak of 30 kW using a battery of 30 kWh. Here, the battery is recharged at the end of the day. Since there often is no energy consumption at night, one can optimize the battery size for each day individually and take the maximum over all days as the final size. For our use case, effects that are only visible for measurement frequencies above 1 Hz are not relevant. Therefore, we consider measurements spaced at a few seconds. Higher frequencies would for example be useful for systems for real-time charging decisions.

In a nutshell, peak shaving depends on two factors. On the one hand, the consumption determines the potential for peak shaving. Naturally, a flat consumption profile has less potential than one with a steep, short peak in the middle of the day. In our case, the profiles are given as a history of smart-meter measurements. On the other hand, feasibility depends on the battery system configuration, i.e., the maximum capacity, the maximum power, the efficiency, and the strategy for charging and discharging. We formalize this as the optimization problem MinBat, see Figure 2. It builds upon a Linear Program that has been proposed to minimize the cost of PV and storage installations [6].

[1]Our implementation and data is publicly available at https://github.com/energystatusdata/aggregation-peak-shaving

$$
\begin{array}{lllr}
\text{MinBat:} & \underset{P_c, P_d}{\text{minimize}}\ C & & \\
\text{subject to} & b(1) = C, & & \text{(a)} \\
& b(T) = C, & & \text{(b)} \\
& b(t) \leq C, & t = 1, \ldots, T & \text{(c)} \\
& b(t) = b(t-1) & & \text{(d)} \\
& \quad + P_c(t-1) \cdot \eta \cdot \Delta t & & \\
& \quad - P_d(t-1) \cdot (1/\eta) \cdot \Delta t, & t = 2, \ldots, T & \\
& P_c(T) = 0, & & \text{(e)} \\
& P_d(T) = 0, & & \text{(f)} \\
& P_c(t) \cdot \eta \cdot \Delta t \leq C - b(t), & t = 1, \ldots, T & \text{(g)} \\
& P_d(t) \cdot (1/\eta) \cdot \Delta t \leq b(t), & t = 1, \ldots, T & \text{(h)} \\
& P_c(t) \leq \max(\hat{L} - L(t), 0) \cdot \eta, & t = 1, \ldots, T & \text{(i)} \\
& P_d(t) = \max(L(t) - \hat{L}, 0). & t = 1, \ldots, T & \text{(j)}
\end{array}
$$

Figure 2: Linear program to minimize battery capacity C.

Objective. The objective is to minimize battery capacity C by deciding on how much to charge ($P_c(t)$) or discharge *power* ($P_d(t)$) between t and $t+1$. This is, $P_c(t) \cdot \Delta t$ and $P_d(t) \cdot \Delta t$ is the *energy* the battery is charged or discharged with.

Parameters. The optimization problem is parametrized by the consumption $L(t)$ (in kW), the target maximum peak $\hat{L}$, the time between consumption measurements Δt, and the maximum time step T. The battery efficiency η accounts for losses when charging and discharging the battery.

Constraints. The state of charge of the battery $b(t)$ (in kWh) is full at the beginning of the time series (a) and at the end of the time series (b). The charge of the battery is restricted by its capacity C (c). The battery charge increases (decreases) during each time step of length Δt by the charging (discharging) power, multiplied with Δt (d). At T, the battery is neither charged nor discharged (e and f). The maximum charge and discharge of the battery depend on its state of charge (g and h). Charging the battery must not exceed $\hat{L}$ (i). The battery compensates loads that surpass $\hat{L}$ (j).

Assumptions. In our optimization, we make two assumptions. First, we choose a fixed battery efficiency η. The round-trip efficiency η_{batt} of lithium-ion batteries is between 95 % and 98 % [5]. With equal losses during charging and discharging, $\eta = \sqrt{\eta_{\text{batt}}}$ is between 97 % and 99 %. In the following, we assume $\eta = 97$ %. Second, we do not consider the maximum charging and discharging power. The maximum power provided by a lithium-ion battery depends on its specific type, the length of the discharge pulse, the cell temperature, and the age. In our use case, the maximum discharge required is 55 kW, which most modern high-power cells achieve. Thus, we do not consider the maximum power in the constraints.

Aggregation. To study the effect of data aggregation, we aggregate the consumption measurements $L(t)$ over time. We use aggregation over non-overlapping windows of length Δt. We consider

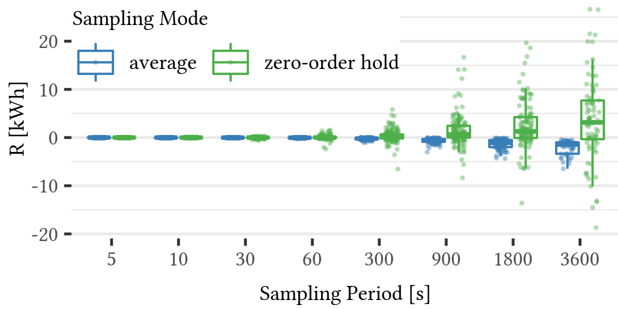

Figure 3: The error for different aggregation intervals for a target maximum peak of 40 kW.

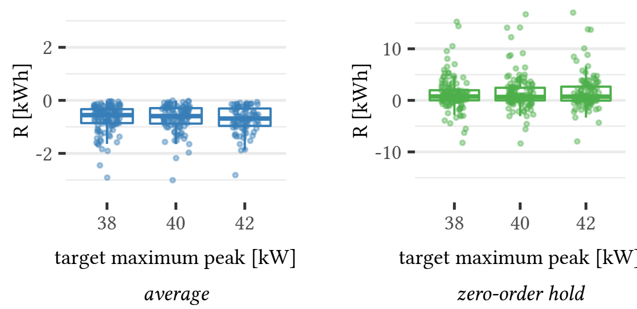

Figure 4: The error for different target maximum peaks for an aggregation interval of 900 s.

two types of aggregation. The first one is *averaging*, i.e., taking the average value over each time window. The second one is *zero-order hold* (ZOH) sampling. For each window, we take the first value. We call Δt, i.e., the number of seconds per window, an *aggregation interval*. For a day worth of measurements, $\Delta t = 1$ h results in 24 aggregated values, and $\Delta t = 1$ min in $24 \cdot 60 = 1440$ values.

3 EXPERIMENTS

Our experiment consists of two parts. First, we evaluate the effects of aggregation on the result of battery size optimization. Secondly, we use data mining to correct the battery size obtained from aggregated data based on a few days of fine-granular measurements.

3.1 Setup

In our experiments, we use high-resolution consumption data from a manufacturing site for electronic components [3] as a ground truth. In particular, we use data from the main terminal, i.e., the overall consumption of the site. The data set contains 281 days with a varying measurement resolution of a few seconds. The overall peak load is 93.41 kW.

We space the measurements equally to 5 s with linear interpolation. The aggregation intervals in our experiments are $\Delta t \in \{10\,\text{s}, 30\,\text{s}, 1\,\text{min}, 5\,\text{min}, 15\,\text{min}, 30\,\text{min}, 60\,\text{min}\}$. We look at different peak shaving targets $\hat{L} \in \{38\,\text{kW}, 40\,\text{kW}, 42\,\text{kW}, 45\,\text{kW}, 50\,\text{kW}\}$. With a target maximum peak $\hat{L} < 38$ kW the battery cannot be recharged by the end of the day in some cases. This results in non-feasible optimization problems.

3.2 Effects of Aggregation

We determine the ground truth battery size C^* for a single day from the 5 s interpolated time series. The battery size based on aggregated data is $\bar{C}^*$. The error is $R = \bar{C}^* - C^*$. Figure 3 shows the error for different aggregation intervals for a target maximum peak of 40 kW with a maximum capacity $C^*_{\max} = 94.55$ kWh over all days. We observe small errors, i.e., less than 10 kWh, for aggregation intervals of up to 1 min. The error grows with aggregation intervals beyond 15 min up to around 25 kWh. A further observation is that the error size depends on the aggregation type: ZOH results in larger errors than averaging. An explanation is that averaging uses more information since it is a summary of all measurements, while ZOH relies on only one measurement. With averaging, the algorithm tends to underestimate the capacity, while ZOH leads to both over- and underestimates. This is intuitive, since the average is always less than or equal to the maximum in the aggregation window, i.e., it underestimates the peak load.

Next, we compare the error over different target maximum peaks for the aggregation interval 900 s, see Figure 4. For ZOH, results are similar across different target maximum peaks. For averaging, the absolute error tends to become larger for larger target peaks.

Table 1 summarizes $\bar{C}^*_{\max}$, the erroneous maximum capacity over all days, for different aggregation intervals Δt and aggregation types for the two target maximum peaks 38 kWh and 40 kWh. Apart from one case ($\hat{L} = 40$ kWh, $\Delta t = 1800$ s, ZOH), $\bar{C}^*_{\max}$ is less than $C^*_{\max}$, resulting in a negative error, i.e., a size underestimation.

This finding has implications on the decision of the plant manager. A battery that is too large can still provide the peak shaving required. A smaller and cheaper battery would have been sufficient in this case. However, if the battery capacity is too small, the target maximum peak will be surpassed, which in turn results in higher costs for peak consumption. If these costs outweigh the cost of additional battery capacity, one may add a safety margin on the battery size. For our data, a margin of about 30 % is sufficient for aggregation intervals between 5 min and 30 min.

Table 1: $\bar{C}^*_{\max}$ and the error $R_{\max} = \bar{C}^*_{\max} - C^*_{\max}$ for different target maximum peaks $\hat{L}$ and aggregation intervals types.

			zero-order hold		average	
$\hat{L}$ [kW]	$C^*_{\max}$ [kWh]	Δt [s]	$\bar{C}^*_{\max}$ [kWh]	$R_{\max}$ [kWh]	$\bar{C}^*_{\max}$ [kWh]	$R_{\max}$ [kWh]
38	94.55	300	87.99	−6.56	76.26	−18.30
		900	72.90	−21.65	75.94	−18.62
		1800	69.76	−24.79	75.78	−18.77
		3600	55.42	−39.12	61.38	−33.18
40	78.89	300	72.37	−6.52	78.87	−0.03
		900	76.89	−2.00	78.47	−0.43
		1800	81.01	2.12	78.62	−0.28
		3600	57.14	−21.75	59.45	−19.45

Table 2: Improvement $\gamma_{\mathbf{max}}$ for different target maximum peaks $\hat{L}$ and aggregation intervals and types.

		zero-order hold		average	
$\hat{L}$ [kW]	Δt [s]	max γ [kWh]	median γ [kWh]	max γ [kWh]	median γ [kWh]
38	300	0.75	−0.30	0.16	0.10
	900	3.09	0.92	0.40	0.16
	1800	6.48	1.61	0.23	−0.31
	3600	2.15	−4.11	2.73	0.99
40	300	0.65	−0.52	0.21	−0.22
	900	3.50	−2.31	0.41	0.16
	1800	0.73	−2.32	0.99	−0.87
	3600	5.65	2.36	0.36	−0.84

3.3 Predicting Errors on Historic Data

Naturally, adding a blanket safety margin on the resulting capacity can result in overestimating $C^*_{\max}$. The unused additional capacity then increases the cost of the battery. In this section, we use machine learning to predict how much one overestimates the battery capacity estimated from aggregated data.

We consider the following scenario, similar to Example 1. For a given $\hat{L}$, we want to find out from a short set of high frequency data how we need to correct $\bar{C}^*_{\max}$ on a larger data set with a known aggregation interval. So the response variable of our model is the error R. As model input, we derive several features: the 0.1-, 0.25-, 0.5-, 0.75-, and 0.9-quantiles, minimum, maximum, mean, standard deviation, variance and Fourier coefficients of the aggregated measurements. The Fourier coefficients are derived as follows. We first create the single-sided amplitude spectrum of the discrete Fourier transform. We then take 12 windows of equal width in the resulting frequency series and take one average value for each window as a feature. In other words, we derive 12 characteristic values from the frequency spectrum. All in all, 24 values describing the frequency spectrum of the time series are available to the model. In addition, we provide the model with $\bar{C}^*_{\max}$, i.e., the estimated capacity on aggregated data.

We take 70 days from the data, which is about 25 % of the days as the training data. The rest of the days are test data. On the training data, we train a stochastic gradient boosting model [4]. To assess the quality of the method, we create 20 such scenarios by randomly splitting the data as described above. For each scenario, we determine $C^*_{\max}$ and $\bar{C}^*_{\max}$ for the days in the test data set. Lastly, we use our model to predict R for each day and add it to $\bar{C}^*_{\max}$. We call this corrected value $\bar{C}^{*+}$.

Over all days that have been corrected this way, we again calculate the maximum $\bar{C}^{*+}_{\max}$. We define the overall improvement through machine learning as the reduction of the error $\gamma = |C^*_{\max} - \bar{C}^*_{\max}| - |C^*_{\max} - C^{*+}_{\max}|$. Table 2 displays statistical data about all γ for different $\hat{L}$ and Δt. We see that the models cannot be used to reliably increase the quality of our results. However, there exist cases in which the models could significantly improve the results. This is indicated by a maximum γ which is larger than 3 kWh in 4 of the observed cases.

4 CONCLUSIONS

We have studied the relationship between the quality of energy-consumption data and the one of peak shaving derived from it. Specifically, we have analyzed how this optimization problem reacts to consumption data that is aggregated with different aggregation intervals and types. In our experiments, aggregation intervals that are larger than 5 min lead to significant errors in the sizing decision for a set of days. Adding a safety margin to the battery size yields a battery capacity that is capable of shaving to the desired maximum peak load. However, this comes at the risk of adding too much unused capacity to the battery. To approximate the real values more precisely, we have trained machine learning models on the resulting error of aggregation. In our experiments, models trained on a short period of high-resolution data do not reliably reduce this error. Currently, adding a safety margin obtained from historic data to the battery size is a more reliable way of reducing the error.

ACKNOWLEDGMENTS

This work was funded by the German Research Foundation (DFG) as part of the Research Training Group GRK 2153: *Energy Status Data – Informatics Methods for its Collection, Analysis and Exploitation.*

REFERENCES

[1] Kumuthawathe Ananda-Rao, Rosnazri Ali, and S. Taniselass. 2017. Battery energy storage system assessment in a designed battery controller for load leveling and peak shaving applications. *Journal of Renewable and Sustainable Energy* 9, 4 (2017), 044107. https://doi.org/10.1063/1.4991455

[2] Lukas Barth, Veit Hagenmeyer, Nicole Ludwig, and Dorothea Wagner. 2018. How Much Demand Side Flexibility Do We Need?: Analyzing Where to Exploit Flexibility in Industrial Processes. In *Proceedings of the Ninth International Conference on Future Energy Systems (e-Energy '18).* ACM, New York, NY, USA, 43–62. https://doi.org/10.1145/3208903.3208909

[3] Simon Bischof, Holger Trittenbach, Michael Vollmer, Dominik Werle, Thomas Blank, and Klemens Böhm. 2018. HIPE – an Energy-Status-Data Set from Industrial Production. In *Proceedings of the Ninth International Conference on Future Energy Systems (e-Energy 2018).* ACM, New York, NY, USA, 599–603. https://doi.org/10.1145/3208903.3210278

[4] Jerome H. Friedman. 2002. Stochastic gradient boosting. *Computational Statistics & Data Analysis* 38, 4 (2002), 367–378. https://doi.org/10.1016/S0167-9473(01)00065-2

[5] Holger C. Hesse, Rodrigo Martins, Petr Musilek, Maik Naumann, Cong Nam Truong, and Andreas Jossen. 2017. Economic Optimization of Component Sizing for Residential Battery Storage Systems. *Energies* 10, 7, Article 835 (2017). https://doi.org/10.3390/en10070835

[6] Fiodar Kazhamiaka, Yashar Ghiassi-Farrokhfal, Srinivasan Keshav, and Catherine Rosenberg. 2018. Robust and Practical Approaches for Solar PV and Storage Sizing. In *Proceedings of the Ninth International Conference on Future Energy Systems (e-Energy 2018).* ACM, 146–156. https://doi.org/10.1145/3208903.3208935

[7] Peter D. Lund, Juuso Lindgren, Jani Mikkola, and Jyri Salpakari. 2015. Review of energy system flexibility measures to enable high levels of variable renewable electricity. *Renewable and Sustainable Energy Reviews* 45 (2015), 785–807. https://doi.org/10.1016/j.rser.2015.01.057

[8] Xing Luo, Jihong Wang, Mark Dooner, and Jonathan Clarke. 2015. Overview of current development in electrical energy storage technologies and the application potential in power system operation. *Applied Energy* 137 (2015), 511–536. https://doi.org/10.1016/j.apenergy.2014.09.081

[9] Alexandre Oudalov, Rachid Cherkaoui, and Antoine Beguin. 2007. Sizing and Optimal Operation of Battery Energy Storage System for Peak Shaving Application. In *Proceedings of the 2007 IEEE Lausanne Power Tech.* IEEE, 621–625. https://doi.org/10.1109/PCT.2007.4538388

[10] Holger Trittenbach, Jakob Bach, and Klemens Böhm. 2018. On the Tradeoff between Energy Data Aggregation and Clustering Quality. In *Proceedings of the Ninth International Conference on Future Energy Systems (e-Energy 2018).* ACM, New York, NY, USA, 399–401. https://doi.org/10.1145/3208903.3212038